

# The Future of Multi-Level Secure (MLS) Information Systems

Chair: John Campbell, Intelink Management Office

Panelists: Tim Ehrsam, Oracle Corporation  
Mario Tinto, The Aerospace Corporation  
Jim Williams, Booz-Allen & Hamilton Inc.

What is the future of MLS Systems? Traditionally, they were thought of as the most secure of systems, occupying the upper levels (classes B1 through A1) of the Department of Defense Trusted Computer System Evaluation Criteria, otherwise known as the "Orange Book". Implementations, to date, have been limited. Today, more people realize that security is a critical part of information systems. So:

- What will the role or future of MLS Systems be?
- In fact, what is a Multi-Level Secure (MLS) Information System? Has the definition remained the same in today's widely connected web-enabled world?
- Are vendors keeping their MLS versions current?
- Why build and use a MLS system?
- What system examples exist?
- What tools exist?
- Can you build successful MLS systems?
- How trusted are the components of the system?
- How do you mitigate lack of trust in the components?
- How do you build such a system?
- How do you measure the security strengths of the system?
- Who are using them?
- What types of systems are being built?
- What effect will size have on MLS systems?
- What measures of security exist to assure that you have a secure system (Standards, criteria, techniques, and methodologies)?
- How do you certify and accredit?

## Some of the Chair's thoughts:

Issue: What is a Multi-Level Secure Information System?

The Orange Book: defines "Multilevel Secure" as a class of systems containing information with different sensitivities that simultaneously permits access by users with different security without risk of compromise.

- Jim Williams pointed out that the DISA Home Page defines Multi-Level Security as: "
- A capability that allows information about different sensitivities (classifications) to be stored in an information system.
  - Allows users having different clearances, authorizations, and need to know the ability to process information in the same system.
  - Prevents users from accessing information for which they are not cleared, do not have authorization, or do not have a need to know "

To me, a Multilevel-Secure Information System is one that satisfies either of the above definitions. The definitions remain the same; implementations may be different.

Issue: Why build and use a MLS System?

The "Why" one uses multilevel security here is frequently due to the "ease-of-use" and data integrity considerations. For example, the user may not want multiple copies of the data. Which copy would be correct? How would you keep all the versions of the same data in synch? And if the same piece of information, say a social security number, has different values, which is the correct value?

A second example is to ease the burden and increase reliability by using of MLS systems, such as workstations and guards, to downgrade data.

This downgrading may be done for several, or more groups of users. The data is stored in separate bins, from which appropriate group can extract the data. MDDS, of which you will be hearing about later, uses this technique.

A fourth example is a distributed database demonstration that we produced. Users at each security level received appropriate multimedia updates, in near real time, something that would have been difficult without this architecture. Each user had all the data that was appropriate for him to have.

"Why" is also answered in an attempt to avoid multiple copies of equipment, networks, etc..

Issue: Can you build successful MLS systems?

Yes, if you know system requirements, threats, vulnerabilities, environment, and use correct methodologies. Protection can be onion-skinned. The risks may be greater than you are willing to accept. If so, requirements, environment and/or requirements may have to be changed and mitigating safeguards may have to be added.

I believe that the concerns about building a MLS system are similar to those about building enterprise-level data warehouses and that the solutions are similar. A belief has been that enterprise-level data warehouses are too big and difficult to build. Such projects are difficult, and some have failed. But with the right team and the right approach, failure is not inevitable. The problem is one of skills, project management and other organizational issues that take time to get right. The people arguing for a quick solution are often those with hardware and software to sell and minimum knowledge or interest in project-management issues. (Barry Devlin, p. 11, DB2 Magazine, summer, 1998.) Again, I believe that the same problems, attitudes and issues exist in the building of multilevel systems.

Issue: What effect will size have on MLS systems?

Systems are getting larger and are containing much more data and the type of data, such as images and other multimedia, takes more space,. For example, database systems are getting bigger. Examples include United Parcel with 11 Terabytes in one database, and 324 billion rows in a DB2 database. (Joe Celko, p. 24, DBMS, June 1998). With the larger systems, even with cheaper storage, there will be more pressure to avoid needless repetition of data and therefore more attempts at MLS systems.

Issue: How will these systems be certified and accredited?

Today, in many cases, different groups own their own information, and share it under predetermined rules, over a system. These systems, called Federated Systems, were a research topic in database security a few

years ago, but are very real now. One example of a Federated System is Intelink, where hundreds of independent groups provide data. In a system such as this, no one person, by himself, should certify the addition of a new system to the federation. The security of the entire system is only as strong as its weakest link, and, adding a weak link may put an existing user's data at risk. Therefore, teams of security information specialists, representing the data owners, should be certifying new systems. This is the approach now being taken in Intelink.

Issue: How do you build an MLS system?

We are fortunate to have additional tools available to us today, which, when used properly, increase the security in systems. Public key and encryption technologies can be used to provide strong identification and authentication, preserve integrity and improve separation and secrecy. These tools should be considered in the architectures of future systems.

## The Panel:

We have three panelists that work on the leading edge of system security. The panelists have differing backgrounds so that they can provide differing perspectives and insights into these issues.

Tim Ehram is a Manager of Security Products at the Oracle Corporation. Oracle is a leading vendor of trusted database products and other products that support security in database systems. He has over fifteen years of experience in trusted system design, development and marketing, the last nine of them being at Oracle. In addition to using this extensive experience to address items from the above list of questions, he will look at the problem from the vendor's point of view. What techniques, hardware and software are available to develop, build and field a multilevel system? Which are being used and who are using them? What are the results of their efforts?

Mario Tinto is an Engineering Specialist in the Trusted Computer Systems Department at the Aerospace Corporation. Prior to this position, he had an outstanding career at the NSA, a leader in the Common Criteria development, a lead evaluator and standards developer and a technical resource that was used to solve hard problems. In addition to addressing items from the above list of questions, he will discuss relevant security standards. Should the Common Criteria be used to develop a secure MLS system? Are there any other standards that would be useful? How do we get "warm fuzzies" about the security of a system? When is a system sufficiently secure for its intended purpose?

Jim Williams has recently become a technical member of Booz-Allen & Hamilton. Prior to this he was a Lieutenant Commander, USN, in the Architecture & Systems Division of USCINCPAC/J21. In both positions, he is/was a lead in the development of the Multi-Domain Dissemination System (MDDS), a very interesting system and a system that if successful, will be widely imitated in that user's community. The system is a subsystem in a much larger system. As such, the subsystem must maintain a level of security that not only protects itself, but also protects access to the larger system through the subsystem. Jim has also been a lead in the installation and use of other multilevel capabilities in CINCPAC. In addition to addressing items from the above list of questions, Jim will describe some of his MLS experiences. Which MLS devices does his employer use? What are their experiences? What could be improved in future systems? What is the MDDS? Why are other users interested in it? What has he learned in building the MDDS?

John Campbell is currently serving in the Security Policy and Plans Directorate in the Intelink Management Office. Here he develops plans, policies, prototypes and solutions to make the very successful Intelink more successful and more secure. Prior to that he has served 13 years at the NSA. He worked with vendors to develop secure information systems and trusted components. He also worked on standards such as the Trusted Database Interpretation and unified INFOSEC and ANSI Criteria, and championed the technical

health of NSA through service on career, technical track and inter-organization panels. He has also served as security lead, system developer, contract manager, professor, programmer, systems analyst, systems security analyst, statistician, mathematician, actuary and operations research analyst. He holds five degrees.

## **New Directions for Multilevel Secure Information Systems**

Tim Ehram, Oracle Corporation

As we examine the future of multilevel secure (MLS) systems, images of the Black Knight in Monty Python's "The Holy Grail" come to mind. In that classic scene, as various appendages of the knight are hacked off, the knight refuses to admit defeat. "It's only a flesh wound!" declares the—knight as his arm is amputated, The market for MLS technology may be in a similar state. The question is, will the "MLS Knight" ignore market realities (just a flesh wound), or can it transform itself into something that can be relied upon to defend our information systems?

Before we discuss the future of MLS, let's look at MLS past from one vendor's perspective. In 1993, Oracle Corporation introduced Trusted Oracle7, a multilevel secure database management system (DBMS). Two years later, at a cost of approximately one million dollars (US), Trusted Oracle7 achieved U.S. TCSEC B1 and ITSEC E3 evaluation ratings. Trusted Oracle continues to provide Oracle with sufficient return on investment to warrant continued support and enhancements in order to meet customer demand, in addition to complying with worldwide security evaluations.

However, compliance with existing security standards is often at odds with the timely release of functionality desired by the market. The Trusted Database Interpretation (TDI) applied strict Orange Book B1 criteria to DBMSs, even though the MLS policy model was not immediately appropriate to most non-defense customers. With the "If you build it, we will buy it" promise, commercial DBMS providers used the TDI as the model for designing MLS DBMS product, since the "B1 stamp" was viewed as the only way to gain entry into the "lucrative" MLS DBMS market. Indeed, most international government programs specified B1 or equivalent functionality/assurance products. Yet, few programs actually implemented the MLS technologies they required vendors to provide (even though the proposed products met the program requirements). Instead, they either integrated a suite of non-MLS components to "simulate" MLS functionality, or they continued to use separate "system high" systems. Although it can be argued that the cost of simulating MLS was usually greater than implementing true MLS products, customers continued to purchase non-MLS products. The result was that many providers of MLS technology could not generate enough business to justify resources to support or enhance existing MLS products. MLS technology releases lagged non-MLS technology releases,

Based upon this experience, Oracle reexamined the market, not just the MLS market. With respect to security, the results indicate that we need to look at the functionality and assurance of B1 and MLS in a new way. The challenge for the future is to provide flexible features that support MLS-style access control and sufficient assurance to meet operational requirements and manage risk. While some (very few, it seems) customers require a B1 DBMS on a B1 operating system, many want strong access control mechanisms (similar to B1 Mandatory Access Control) with less rigid policy models than exist on today's MLS systems. This takes the form of higher assurance "information flow control," to enable sites to describe their own policy. Users also require the management tools to easily configure and implement these policies as they apply to their business – be it Defense, healthcare, education, banking, etc. Additionally, they require these mechanisms on ANY operating system, not just B1 operating systems. Naturally, the technology must support the latest and greatest hardware and software.

Traditional MLS technology vendors (those still standing) are beginning to release products that view MLS functionality as an integrated component of their mainstream offering. The ability of MLS vendors to build products that satisfy the diverse operational and assurance requirements of all market sectors will determine whether products that support MLS architectures will be available in the future. The market can no longer deny that we have more than a flesh wound. "Come back here, you ruddy coward!"

# The Future of MLS Information Systems

Mario Tinto  
The Aerospace Corporation

## I. The Fundamental Issues Persist

The problems envisioned during the 70's & 80's are upon us with a vengeance today;

- We have widely distributed systems. It is common for an enterprise to have elements literally around the globe.
- That we are internetted goes without saying; it is not stretching the truth to state that with today's Internet, everyone is connected to everyone else,
- We once worried about the problems of performing authentication across a network connection, "mobile user" primarily referred to one or two military aircraft during periods of national emergency. Today everyone is somewhere else (i.e., remote from the server being accessed), the notion "mobile user" refers to anyone with a cell phone who wants to surf the net while cruising the highways.
- Where we worried about, and created procedural controls for "unauthorized software" (typically introduced via floppies), we now are capable of instantly-and on demand-downloading arbitrary executables from a myriad of web sites; whether games, interesting utilities (e.g., mailers, audio and quickmovie players, browsers), and even operating system fixes and upgrades. Some executables come unbidden; via programming or scripting languages embedded in web pages.

## II. The Need Continues

The need for the capability to protect information and resources is as great as it ever was. Commercial companies host their web sites on servers that are networked with other, if not all, of the organization's computer assets. Government and military agencies need to communicate across the Internet and access the web from workstations that are also networked with systems storing and processing sensitive data. In short, the subversion of a single client or server provides an attacker with immediate connectivity to the information and computing resources of an entire organization.

The concerns that stimulated the growth of the computer security community in the mid 70's were fundamentally the same as those facing us today-the need to protect the operating system and sensitive information from potentially malicious code. However, the scope of the problem that faces us today is considerably greater. It wasn't too many years ago that the primary method for introducing code into a computer system was by physically loading a floppy disk. Today, code may be introduced into a client merely by clicking on a hypertext link.

## III. The Technology is Available

The basic notions and technologies developed during the 1970's and 1980's are still valid and effective. The notion of a computer system that is self-protecting, and is capable of enforcing a set of access control rules is fundamental to obtaining the flexibility and utility we need and desire in our computing systems, while providing protection of vital enterprise assets.

There is sufficient need to justify a real future for MLS systems. Security is not just a "nice-to-have" add-on; it is a vital operational requirement. And, we do not want for the technology. What has been wanting has been our resolve to spend our dollars such as to demonstrate the market demand to which developers would respond. The question then, is only whether we want fundamental solutions, or are going to go into the future with ineffectual patches, partial solutions, and piecemeal approaches to serious security problems.